SINk: A Middleware for Synchronization of Heterogeneous Software Interfaces

Mohammad Hosseini*, Yu Jiang*, Poliang Wu*, Richard B. Berlin Jr.*†, Lui Sha*

*Department of Computer Science
University of Illinois at Urbana-Champaign (UIUC)

†Department of Surgery
Carle Foundation Hospital

{shossen2, jy1989, wu87, rberlin, rberlin, lrs}@illinois.edu

arXiv:1510.05606v1 [cs.SE] 19 Oct 2015

ABSTRACT

Software is everywhere. The increasing requirement of supporting a wide variety of domains has rapidly increased the complexity of software systems, making them hard to maintain and the training process harder for end-users, which in turn ultimately demanded the development of user-friendly application software with simple interfaces that makes them easy, especially for non-professional personnel, to employ, and interact with.

However, due to the lack of source code access for third-party software and the lack of non-graphical interfaces such as web-services or RMI (Remote Method Invocation) access to application functionality, synchronization between heterogeneous closed-box software interfaces and a user-friendly version of those interfaces has become a major challenge. In this paper, we design SINk[1], a middleware that enables synchronization of multiple heterogeneous software applications, using only graphical interface, without the need for source code access or access to the entire platform's control. SINk helps with synchronization of closed-box industry software, where in fact the only possible way of communication is through software interfaces. It leverages efficient client sever architecture, socket based protocol, adaptation to resolution changes, and parameter mapping mechanisms to transfer control events to ensure the real-time requirements of synchronization among multiple interfaces are met. Our proof-of-concept evaluation shows there is in fact potential usage of our middleware in a wide variety of domains.

Categories and Subject Descriptors

H.5.2 [**User Interfaces**]: Graphical user interfaces (GUI)

General Terms

Design, Experimentation

[1]A demo illustrating how our middleware works in practice is available at http://publish.illinois.edu/mdpnp-architecture/?p=639

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ARM 2015, December 07-11, 2015, Vancouver, BC, Canada
Copyright 2015 ACM 978-1-4503-3733-5/15/12 ...$15.00
http://dx.doi.org/10.1145/2834965.2834967.

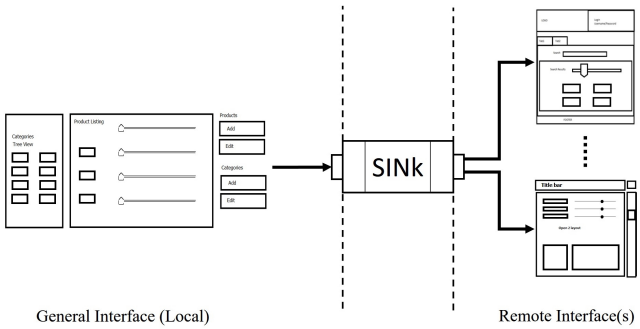

Figure 1: The SINk workflow. Multiple graphical interfaces (right-side) are remotely synchronized with a single interface (left).

1. INTRODUCTION

"Software is eating the world!" [9]. Our dependency on software is continuously increasing, and it is said that 60-90% of production in the automotive domain for example, is done by software systems [21]. Many products and industrial services that would have traditionally been realized through "hardware", are now realized purely via "software solutions". Overall, one way or another, human-in-the-loop software systems in various domains are getting more and more complex, as they operate within a complex ecosystem of libraries, models, protocols and devices, and require human interaction [21]. The interfaces of many platform-dependent software, such as industrial controller simulators (e.g. Mitsubishi PLC x7 [20]) and healthcare systems (e.g. Laerdal's SimMan Patient Simulator [6]) for example, are sometimes hard to manage and lack user-friendliness. Therefore, third-parties are pushed to develop simple-to-use, and more user-friendly and maneuverable interfaces for those applications, which in fact motivates the need for co-simulation among different interfaces. While the graphical user interfaces are easy to develop, there has been a significant demand on interface-to-interface synchronization of heterogeneous software interfaces.

Unfortunately, existing tools such as those spanning from remote desktop applications and desktop sharing to collaborative software applications lack support of interface synchronization, and only provide access to applications simply through desktop screen sharing and manual control by users.

In this paper, we describe SINk, a middleware that enables interface-to-interface synchronization and automatic

remote control of heterogeneous graphical interfaces. The design of this middleware is mainly motivated by connecting and synchronizing heterogeneous applications with homogeneous functionality, but different graphical interfaces, over a network or the Internet, as presented in Figure 1. Once the interfaces are connected through the established client-server connection, our middleware allows users to automatically synchronize multiple applications, remotely, without physical access or visual view of the remote desktops, as if they were sitting right in front of the remote applications. In this manner, we enable the user to control any remote applications and automatically perform actions such as opening and closing windows and tabs, pushing buttons, applying keystrokes, and updating strings, values, and checkboxes, using only their graphical interfaces and without causing mismatches between them. SINk is adaptive in the sense that it can accommodate varying platform features such as changes in display screen resolutions, as it can dynamically adapt itself to locate pixel values *relative* to any resolution. Moreover, SINk's automated mechanisms achieved through *interface-only* control incurs a high degree of flexibility, and can effectively adapt to ecosystem changes when reconfiguration of application rapidly occurs. That leads to significant reduction in heterogeneous software maintenance costs.

Technically speaking, the SINk middleware leverages efficient architecture, protocol, and parameter mapping mechanisms to transfer control events, while at the same time ensuring consistency, bandwidth saving, platform independence and the fulfillment of real-time requirements for synchronization. In summary, SINk

- automatically performs remote control as opposed to manual control by users,
- does not require visual view of remote desktop, thus providing significant bandwidth savings,
- does not require source code or non-graphical interfaces (such as web services or RMI) access to remote applications,
- performs synchronization in real-time,
- is platform independent.

To the best of our knowledge, no single middleware currently exists that achieves synchronization among heterogeneous applications in a coherent way. Moreover, SINk can further assist software engineers to build a single user-friendly interface as a general application interface for front-end interaction, or to realize co-simulation among multiple heterogeneous graphical interfaces.

2. RELATED WORK

SINk is conceptually similar to the notion of *mediators* underlying emergent connectors [16, 17, 11] such as Enterprise Service Bus [10] as the concept of a "connectivity middleware" is common between the two. However, SINk is fundamentally different as the design goal of mediators is to enable the composition of pervasive networked systems, protocol mediation, and interoperability in distributed systems as opposed to remote interface-based synchronization. The most related tools to SINk are remote desktop and desktop sharing software, which allow a personal computer's desktop environment to be run remotely on one system, while being displayed on a separate client device. Microsoft's Remote Desktop Connection [5], Apple Remote Desktop [1], and Chrome Remote Desktop [3] for example, allow users to remotely connect to a computer from another computer, therefore providing access to programs and files by visually controlling the keyboard and mouse and relaying the graphical screen over a network. Similarly, desktop sharing applications and collaborative software such as Microsoft's Lync [4] and TeamViewer [7] provide desktop access to a remote machine running the same software, helping users to remotely control and share a desktop, with the additional option of video conferencing services. However, not only are these applications manual and user-controlled, but the remote desktop software and desktop sharing applications also act in a *computer-to-computer* manner and have computer-wide access. This is not easily applied to the synchronization among platform dependent applications, such as flight control and autopilot systems in drones [12], automatic remotely-controlled construction machinery in smart-grids [22], and co-simulation of heterogeneous production and ERP software in the automotive industry [8].

In SINk , on the contrary, the notion of access is platform independent and lightweight because it is *application-to-application* or *interface-to-interface*. Furthermore, control and input parameters are directed *automatically* into the remote graphical application interfaces residing on the remote computers, thus synchronizing multiple interfaces and allowing users to need only control a *single* interface. In addition, SINk eliminates the unnecessary need to visually share the desktop views, hence allowing for significant bandwidth savings by avoiding the real-time encoding and transfer of desktop views, especially crucial for power-limited mobile devices [13, 15]. Moreover, remote users have no ability to modify the shared content and resources whatsoever, and are only passively controlling remote interfaces.

3. DESIGN OF THE MIDDLEWARE

SINk is implemented through a mapping system as well as a communication system accomplished through a client/server architecture. The client is installed on the local computer running the local application and then connects to the server component, which is installed on the remote computer. During SINk sessions, all corresponding keystrokes and mouse clicks are registered as if the users were actually sitting in front of the remote computers and performing tasks on the remote interfaces. We implemented SINk in Java that can be deployed on any platform running Java Virtual Machine (JVM), including Linux and Windows. Therefore, JVM is a base requirement, making the compiled code platform-independent. We have designed a list of APIs for the users, such as performing a remote connection, specifying control attributes, and transferring parameter values.

3.1 Middleware Structure

SINk consists of three major components: a local agent (control client) residing on the local machine, a mapping module, and a remote agent (control server) residing on each remote machine, as illustrated in Figure 2. The local agent communicates with the remote agents through a secure persistent message exchange communication system. Users' control inputs are received by the local agent, encoded to a specific message format, and are then directed to the mapping module.

The mapping module is pre-configured with interface control attributes to provide a particular set of interface func-

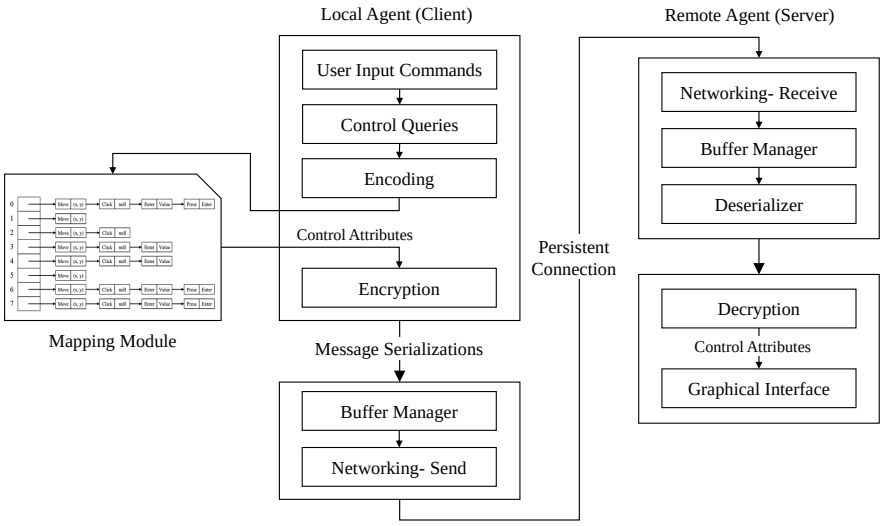

Figure 2: The overall structure of SINk.

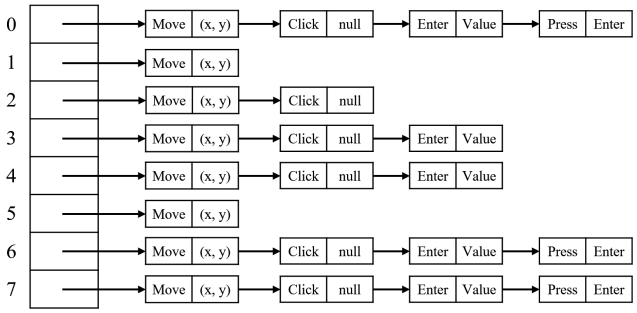

Figure 3: An example mapping module.

tions on each of the remote interfaces. This happens by performing transformation of local control inputs on the local platform to remote interface control attributes corresponding to the remote graphical interfaces on the remote platform, thereby allowing remote synchronization with each remote interface. This is similar to a remote desktop connection, but it happens automatically and with no need for sharing the desktop view or visual control with users. The output data originating from the mapping module consists of control attributes, which are then encrypted with the AES 128-bit symmetric cipher in electronic codebook (ECB) mode, buffered, serialized, and then transferred to the remote agent via the persistent socket connection. While placement of the mapping module as a centralized module on the local machine is more convenient for updates, auditing, and security reasons, it is not yet a hard requirement. The module can be placed separately on each individual remote machine alternatively.

The control messages are deserialized and decrypted once received at the remote agent. Remote synchronization between the local interface and remote interfaces is performed via interface control functions in accordance with the control attributes received through the communication channel. Although currently synchronization is only one-way (from local to the remote interfaces), without loss of generality, SINk can be reconfigured so that changes and results on remote interfaces be synchronized back and displayed on the local interface for any possible adaptation purposes.

3.2 Customized Client-Sever Architecture

From an engineering point of view, unlike a regular client-server connection such as those in chat systems with the client looping to read the responses, our middleware tool must also support sporadic message transfer but with no connection termination. However, it also needs to maintain a live and permanent connection after each transfer in order to incur minimum latency.

To address the requirement above, we customized a low-overhead persistent client-server connection over TCP/IP throughout the running session rather than setting up a new connection for each transfer. This maintains the stability of the socket connections by initially creating a connection at the beginning of each session, and occasionally sending a message given the system's input. To enable that, we wrapped the client socket connection around a thread, and use a blocking queue to wait for messages. A single sender queue exists throughout the application, therefore using a singleton pattern. On the other side, performing a `read()` function causes the thread to block forever. To address that, we use a special type of thread that calls a specific method repetitively at specified periods and read time-out that can be used to post a message, a ping message, every so often, which improves the stability of connections while also relaxing problems associated with closing the applications due to calling the `close()` function.

3.3 Data Structures and Rules for Mapping

The mapping module works on the principle of key-value store and hashing, composed of a combination of hash-map and 2-dimensional linked-list data structures, which is used to simulate user interaction and control the graphical interfaces pre-configured with mouse and keyboard events. To store key-value pairs, we used the first dimension of the 2D linked-list as a bucket to store key objects corresponding to

encoded user inputs, while the second dimension is used to store values, corresponding to an ordered list of interface control attributes such as necessary mouse clicks and key presses that must be executed on the remote interfaces to perform identical actions. Similar to a regular *HashMap*, the mapping module's `get(Key k)` method calls *hashCode* method on the key inputs, and applies returned *hashValue* to its own static hash function to find a bucket location where keys and values are stored. Figure 3 shows an overview of an example mapping module. For example, the first entry of the map as shown in Figure 3, executes the following chain of events:

1. *Move* the mouse pointer to a specific 2D coordinate on the display screen (given x and y coordinates as the horizontal and vertical addresses of any pixel, respectively),

2. Perform a mouse *click* event on current pointer (we implemented click event as a combination of mouse left button's *press* and *release* events, with an intermediate delay of 200 ms),

3. *Enter* a specific value or number in the current position. This requires it to iteratively *press* multiple specific keys on the keyboard, and

4. *Press* the "Enter" key on the keyboard.

Our implementation of the mapping module imposes a one-time overhead, and it can be reusable. Therefore, if the application's user interface changes considerably, only the mapping module is updated, incurring minimum cost so the automation does not need to be rewritten.

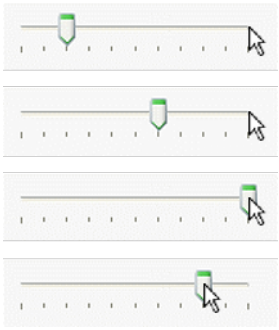

Figure 4: Interface control process for an example horizontal slider bar.

The graphical interface control attributes are implemented as a series of sequential mouse and keyboard events. While many of the graphical user interface components are elementary actions and are straightforward to control through events originating from multiple registered mice and keyboards, such as moving the mouse pointer to a specific coordinate location, clicking, pressing a key or entering a value, interesting challenges exist when controlling or adjusting some interface components such as scroll bars and slider bars. Let's take a horizontal slider bar for example. To control a slider bar to set a new value, use of a mouse drag event is infeasible as the initial position of the slider knob is unknown for the mouse pointer to hover on. While the current position of the slider knob is not known, the coordinates of minimum or maximum value endpoints are known on the horizontal bar. The design trick to address this challenging issue is to first move the mouse pointer to the coordinate corresponding to either endpoint, and then perform mouse clicks multiple times to push the slider pointer on the track towards the specific endpoint (the maximum number of mouse clicks is deterministic- in our experiments the number was four). Once the mouse pointer is on the slider knob, a mouse drag event is then performed to move the knob to the desired pre-determined position corresponding to the desired control value. Figure 4 (top to bottom) visually illustrates the control process. Although graphical component functions and views are subject to operating system, design language and layout variants, mouse and keyboard events can be registered for our control mapping purposes without loss of generality.

4. EVALUATION

We have evaluated and tested SINk rigorously over our industrial case study conducted in collaboration with Carle Foundation Hospital [2], on a real platform where 138 synchronization requirements were specified to synchronize two medical simulator software products. The requirements were inspected multiple times with developers, researchers, and physicians to ensure that specific functional requirements are satisfied.

The evaluation platform is presented in Figure 5, with the closed-box simulator software as the remote interface on the bottom, and the local interface illustrated on the top. The closed-box software is SimMan's [6] advanced patient simulator shipped with a laptop running Windows XP, which controls a SimMan medical manikin used for basic and advanced life Support skill assessment. The SimMan's simulator software allows observation, recognition, and modification of most vital signs which are used in emergency medicine, fed directly to the manikin itself as well as a patient monitor. The local interface likewise, is a patient simulator locally developed for nurses and physicians as a part of a best practice medical system to perform the most relevant medical interventions according to the medical guidelines and protocols. The local patient simulator features a simple, user-friendly, and easily-operated graphical interface with straightforward and uncomplicated control functions to help nurses and physicians avoid complications of using the SimMan's patient simulator. As an example, the local interface incurs a *single* step including 10 parameters to modify the running heart rhythm, whereas the SimMan's patient simulator involves 9 steps, requiring the user to audit 57 different parameters. With SINk, the input values are only controlled through a single user-friendly interface, and are automatically synchronized with those corresponding to the SimMan's simulator, thereby relieving the users from confusing complications and removing the need to double-enter the input values on a second interface. All 138 synchronization requirements are accomplished correctly.

Apart from the case study and the important benefits resulting from using SINk, our middleware was specially regarded for its automation role. Prior to applying our middleware, a technician was hired to replicate, and manually perform the control functions on the SimMan's simulator as a way to synchronize the user-friendly patient simulator with the SimMan's simulator. Thenceforth, automatic synchronization was achieved, removing the human from the loop. Overall, we have received positive feedback from the experts using the middleware. The qualitative feedback we have received is promising, suggesting the middleware might be applicable to large sets of requirements and extended to domains such as co-simulation of heterogeneous production and ERP software in the automotive industry [8].

Figure 5: Real platform testing (Middle). Local interface (Top). Remote closed-box application (Bottom).

5. CONCLUSION AND FUTURE WORK

In this paper, we presented SINk, an adaptive middleware tool that performs interface synchronization automatically, remotely, and without physical access or visual view of remote interfaces, as if users were sitting right in front of the remote software. We tested and evaluated SINk on a real platform, and showed that apart from daily personal applications, there are in fact many potential uses of our middleware in industry services that can not be realized by other means.

We are currently working on an interface attribute recorder that can capture and log interface control inputs on local interfaces and be fed directly into the mapping module, to strengthen the automation and the scalability of the middleware. We can also exploit image segmentation and energy-efficient texture recognition techniques to learn type and position of graphical components on software interfaces, especially when aimed at interfaces on power-constrained mobile devices [19, 18, 14]. In the future, we also plan to systematically evaluate SINk using quantitative metrics.

6. ACKNOWLEDGMENTS

This research is funded in part by NSF CNS 13-29886 and in part by Navy N00014-12-1-0046.

7. REFERENCES

[1] Apple - remote desktop. www.apple.com/remotedesktop.

[2] Carle Foundation Hospital. http://www.carle.org.

[3] Chrome remote desktop - chrome web store - google. https://chrome.google.com/webstore/detail/chrome-remote-desktop.

[4] Lync - microsoft office (currently known as skype for business. http://products.office.com/en-us/skype-for-business/online-meetings.

[5] Remote desktop connection - microsoft windows. http://windows.microsoft.com.

[6] Simman patient simulator, laerdal medical. http://www.laerdal.com/doc/86/SimMan.

[7] Teamviewer: remote control, remote access, & online meeting. http://www.teamviewer.com.

[8] Must-have erp features for the automotive industry. Plex Systems, Manufacturing Business Technology, 2014. http://www.mbtmag.com/articles/2014/01/must-have-erp-features-automotive-industry.

[9] M. Andreessen. Why software is eating the world. *Wall Street Journal*, vol 20, August 2011.

[10] D. Georgakopoulos and M. P. Papazoglou. Enterprise service bus. In *Service-Oriented Computing*, pages 1–28. MIT Press, 1 edition, November 2008.

[11] J. Green, P. Protocol conversion. *Communications, IEEE Transactions on*, 34(3):257–268, Mar 1986.

[12] S. Helton. Fukushima daiichi workers clear debris by remote control. *21st Century Wire*, August 2014. http://21stcenturywire.com/2014/08/07/flight-control-boeings-uninterruptible-autopilot-system-drones-remote-hijacking.

[13] M. Hosseini, D. T. Ahmed, and S. Shirmohammadi. Adaptive 3D texture streaming in M3G-based mobile games. In *Proceedings of the 3rd ACM Multimedia Systems Conference*, MMSys '12, 2012.

[14] M. Hosseini, A. Fedorova, J. Peters, and S. Shirmohammadi. Energy-aware adaptations in mobile 3d graphics. In *Proceedings of the 20th ACM International Conference on Multimedia*, MM '12, 2012.

[15] M. Hosseini, J. Peters, and S. Shirmohammadi. Energy-budget-compliant adaptive 3D texture streaming in mobile games. In *Proceedings of the 4th ACM Multimedia Systems Conference*, MMSys '13, 2013.

[16] V. Issarny, A. Bennaceur, and Y.-D. Bromberg. Middleware-layer connector synthesis: Beyond state of the art in middleware interoperability. In M. Bernardo and V. Issarny, editors, *Formal Methods for Eternal Networked Software Systems*, volume 6659 of *Lecture Notes in Computer Science*, pages 217–255. Springer Berlin Heidelberg, 2011.

[17] V. Issarny, B. Steffen, B. Jonsson, G. Blair, P. Grace, M. Kwiatkowska, R. Calinescu, P. Inverardi, M. Tivoli, A. Bertolino, and A. Sabetta. Connect challenges: Towards emergent connectors for eternal networked systems. In *Engineering of Complex Computer Systems, 2009 14th IEEE International Conference on*, pages 154–161, June 2009.

[18] S. Minaee, M. Fotouhi, and B. H. Khalaj. A geometric approach for fully automatic chromosome segmentation. *CoRR*, abs/1112.4164, 2011.

[19] S. Minaee and Y. Wang. Screen content image segmentation using least absolute deviation fitting. *CoRR*, abs/1501.03755, 2015.

[20] C. M. Park, S. M. Bajimaya, S. C. Park, G. N. Wang, J. G. Kwak, K. H. Han, and M. Chang. Development of virtual simulator for visual validation of plc program. In *Computational Intelligence for Modelling, Control and Automation, 2006 and International Conference on Intelligent Agents, Web Technologies and Internet Commerce, International Conference on*, pages 32–32. IEEE, 2006.

[21] A. Trendowicz. Why software effort estimation? In *Software Cost Estimation, Benchmarking, and Risk Assessment*, The Fraunhofer IESE Series on Software and Systems Engineering, pages 3–7. Springer Berlin Heidelberg, 2013.

[22] M. Williams. Fukushima daiichi workers clear debris by remote control. *Computer World*, April 2011. http://www.computerworld.com/article/2507273/computer-hardware/fukushima-daiichi-workers-clear-debris-by-remote-control.html.